

Situational Confidence Assistance for Lifelong Shared Autonomy

Matthew Zurek*
UC Berkeley
matthewzurek@berkeley.edu

Daniel S. Brown
UC Berkeley
dsbrown@berkeley.edu

Andreea Bobu*
UC Berkeley
abobu@berkeley.edu

Anca D. Dragan
UC Berkeley
anca@berkeley.edu

ABSTRACT

Shared autonomy enables robots to infer user intent and assist in accomplishing it. But when the user wants to do a new task that the robot does not know about, shared autonomy will hinder their performance by attempting to assist them with something that is not their intent. Our key idea is that the robot can detect when its repertoire of intents is insufficient to explain the user's input, and give them back control. This then enables the robot to observe unhindered task execution, learn the new intent behind it, and add it to this repertoire. We demonstrate with both a case study and a user study that our proposed method maintains good performance when the human's intent is in the robot's repertoire, outperforms prior shared autonomy approaches when it isn't, and successfully learns new skills, enabling efficient lifelong learning for confidence-based shared autonomy.

KEYWORDS

shared autonomy, learning from demonstrations

ACM Reference Format:

Matthew Zurek, Andreea Bobu, Daniel S. Brown, and Anca D. Dragan. 2018. Situational Confidence Assistance for Lifelong Shared Autonomy. In *Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

Shared autonomy [1, 2, 5, 7–9, 13, 17–19, 24] enables robots to assist human operators to more effectively accomplish what they want. Rather than directly executing the human's control input, a typical framework for shared autonomy has the robot estimate the human's intent and execute controls that help achieve it [8, 15, 20, 23, 24].

These methods perform well when the robot has access to a set of possible intents the human might have, e.g. the objects the

*Equal contribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Woodstock '18, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>




Figure 1: We propose an approach for lifelong shared autonomy that enables a robot to detect when its set of known human intents is insufficient to explain the current human behavior. Rather than trying to assist for the wrong intent, the robot learns from novel teleoperations to learn a model of the new intent, allowing for lifelong confidence-based assistance.

human might want to reach, or the buttons they might want to push [8, 15]. But in reality, users of these systems will inevitably want to perform tasks that are outside the repertoire of known intents—they might want to reach for a goal the robot does not know about, or perform a totally new task, like pouring their cup of water into the sink. This presents a three-fold challenge for shared autonomy. First, the robot will be unable to help with something it does not know about. Second, and perhaps more importantly, it will attempt to assist with whatever wrong intent it infers, interfering with what the user is trying to do and hindering their performance. This happens when the robot plans in expectation [15], and, as our experiments will demonstrate, it happens even when the robot arbitrates the amount of assistance based on its confidence in the most likely goal [8]. Third, the 100th time the person attempts the new task, it will remain just as difficult as the first time.

Our key idea is that robots should detect that the user is trying something new and give them control. This then presents an opportunity for the robot to observe the new executed trajectory, learn the underlying intent that explains it, and add it to its repertoire so that it can infer and assist for it in the future.

To achieve this, we need two ingredients: 1) a way for the robot to detect its repertoire of intents is insufficient, and 2) a representation of intents that enables learning new tasks throughout its lifetime, adding them to its repertoire, and being able to run inference over

them in a unified way with the initial, known intents. For the latter, we use cost functions as a way to unify goals and general skills like pouring into the same representation. This then enables the former: when the human acts too suboptimally for any of the known cost functions, it suggests that the robot does not have the right set of costs. autonomy new intent, which the robot uses to learn a cost function via Inverse Reinforcement Learning (IRL) [10] and add it to its set of intents. This enables lifelong shared autonomy, where the robot helps when it is confident in what the user wants and learns new intents when it detects that the human is doing something novel, so that it can assist with that intent in the future.

2 CONFIDENCE-AWARE SHARED AUTONOMY

We consider a human teleoperating a dexterous robotic manipulator to perform everyday manipulation tasks. The robot's goal is to assist the person in accomplishing their desired skill by augmenting or changing their input. While the robot is armed with a set of predefined skills the person might want to carry out, the human's desired motion might not be captured by any of them. We propose that since the robot might not understand the person's intentions, it should reason about how confident it is in its predictions in order to avoid assisting for the wrong skill.

2.1 Preliminaries

Formally, let $x \in \mathcal{X}$ be the continuous robot state (e.g. joint angles), and $u \in \mathcal{U}$ the continuous robot action (e.g. joint velocity). The user controls their desired robot configuration by providing continuous inputs $a \in \mathcal{A}$ via an interface (e.g. GUI, joystick, keyboard commands, etc). These inputs are mapped to robot actions through a *direct teleoperation* function $\mathcal{T} : \mathcal{A} \rightarrow \mathcal{U}$. Define a person's trajectory up until time t as the sequence $\xi_{0 \rightarrow t} = (x^0, a^0, \dots, x^t, a^t)$.

The robot is equipped with a set of known intents Θ , one of which may represent the user's desired motion. Each intent is parameterized by a cost function C_θ , which may be hand-engineered or learned from demonstrations via IRL [21, 28]. Our shared autonomy system does not know the intent a priori, but infers it from the human's inputs. Given the user's trajectory so far, $\xi_{0 \rightarrow t}$, a common strategy is to predict the user's intent $\theta \in \Theta$, compute the optimal action for moving accordingly, then augment the user's original input with it [8]. However, what if none of the intents match the human's input? We introduce a shared autonomy formalism where the robot reasons about its confidence in its current set of intents' ability to explain the person's input, and uses that confidence to determine whether (and how much) to assist or whether to ask the human to demonstrate the missing intent.

2.2 Intent Inference

To assist the person, the robot has to predict which of its known tasks they are trying to carry out, if any. We assume the Boltzmann nosily-rational decision model for the human teleoperation [3, 27]:

$$P(\xi | \theta, \beta) = \frac{e^{-\beta C_\theta(\xi)}}{\int_{\tilde{\xi}} e^{-\beta C_\theta(\tilde{\xi})} d\tilde{\xi}}, \quad (1)$$

where the person chooses the trajectory ξ proportional to its exponentiated cost C_θ . The parameter $\beta \in [0, \infty)$ controls how much

the robot expects to observe human input consistent with the intent θ . Typically, β is fixed, recovering the Maximum Entropy IRL observation model [28]. Inspired by work on confidence-aware human-robot interaction [4, 11, 12], we instead reinterpret β as a measure of the robot's *situational confidence* in its ability to explain human data, given the known intents Θ , and we show how the robot can estimate it in Sec. 2.3.

Given Eq. (1), if the cost C_θ of intent θ is additive along the trajectory ξ , we have that:

$$P(\xi_{0 \rightarrow t} | \theta, \beta) = e^{-\beta C_\theta(\xi_{0 \rightarrow t})} \frac{\int_{\tilde{\xi}_{t \rightarrow T}} e^{-\beta C_\theta(\tilde{\xi}_{t \rightarrow T})} d\tilde{\xi}_{t \rightarrow T}}{\int_{\tilde{\xi}_{0 \rightarrow T}} e^{-\beta C_\theta(\tilde{\xi}_{0 \rightarrow T})} d\tilde{\xi}_{0 \rightarrow T}}, \quad (2)$$

where T is the duration of the episode. We follow [8] and approximate these integrals via Laplace's method:

$$P(\xi_{0 \rightarrow t} | \theta, \beta) \approx e^{-\beta(C_\theta(\xi_{0 \rightarrow t}) + C_\theta(\xi_{t \rightarrow T}^*) - C_\theta(\xi_{0 \rightarrow T}^*))} \times \sqrt{\left(\frac{\beta}{2\pi}\right)^{tk} \frac{|\nabla^2 C_\theta(\xi_{0 \rightarrow T}^*)|}{|\nabla^2 C_\theta(\xi_{t \rightarrow T}^*)|}}, \quad (3)$$

where k is the action dimensionality, and the trajectories $\xi_{0 \rightarrow T}^*$ and $\xi_{t \rightarrow T}^*$ are optimal with respect to C_θ and can be computed with any off-the-shelf trajectory optimizer¹.

Now, given a tractable way to compute the likelihood of the human input, the robot can obtain a posterior over intents:

$$P(\theta | \xi_{0 \rightarrow t}, \beta) = \frac{P(\xi_{0 \rightarrow t} | \theta, \beta)}{\sum_{\theta' \in \Theta} P(\xi_{0 \rightarrow t} | \theta', \beta)}, \quad (4)$$

assuming $P(\theta | \beta) = P(\theta)$ and a uniform prior over intents.

Prior inference-based shared autonomy work [8, 15] typically assumes $\beta = 1$. We show that the robot should not be restricted by such an assumption and it, in fact, benefits from estimating $\hat{\beta}$ and reinterpreting it as a confidence.

2.3 Confidence Estimation

In the Boltzmann model in Eq. (1), we see that β determines the variance of the distribution over human trajectories. When β is high, the distribution is peaked around those trajectories ξ with the lowest cost C_θ ; in contrast, a low β makes all trajectories equally likely. We can, thus, reinterpret β to take a useful meaning in shared autonomy: given an intent, β controls how well that intent's cost explains the user's input. A high β for an intent θ indicates that the intent's cost explains the input well and is a good candidate for assistance. A low β on all intents suggests that the robot's intent set is insufficient for explaining the person's trajectory.

We can thus estimate β and use it for assistance. Using the likelihood function in Eq. (3), we write the β posterior

$$P(\beta | \xi_{0 \rightarrow t}, \theta) = \frac{P(\xi_{0 \rightarrow t} | \theta, \beta)P(\beta)}{\int_{\tilde{\beta}} P(\xi_{0 \rightarrow t} | \theta, \tilde{\beta})P(\tilde{\beta})d\tilde{\beta}}. \quad (5)$$

If we assume a uniform prior $P(\beta)$, we may compute an estimate of the confidence parameter β per intent θ via a maximum likelihood estimate:

$$\hat{\beta}_\theta = \arg \max_{\tilde{\beta}} e^{-\tilde{\beta}(C_\theta(\xi_{0 \rightarrow t}) + C_\theta(\xi_{t \rightarrow T}^*) - C_\theta(\xi_{0 \rightarrow T}^*))} \left(\frac{\tilde{\beta}}{2\pi}\right)^{\frac{tk}{2}}, \quad (6)$$

¹We use TrajOpt [26], based on sequential quadratic programming.

where we drop the Hessians since they don't depend on β . Setting the derivative of the objective in Eq. (6) to zero and solving for β yields the following estimate:

$$\hat{\beta}_\theta^{MLE} = \frac{tk}{2(C_\theta(\xi_{0 \rightarrow t}) + C_\theta(\xi_{t \rightarrow T}^*) - C_\theta(\xi_{0 \rightarrow T}^*))}. \quad (7)$$

The denominator in Equations 7 can be interpreted as the “suboptimality” of the observed partial trajectory $\xi_{0 \rightarrow t}$ compared to the cost of the optimal trajectory for the particular θ , $C_\theta(\xi_{0 \rightarrow T}^*)$. Note that $\hat{\beta}_\theta$ is inversely proportional to the suboptimality divided by the number of time steps t that have passed. If this normalized suboptimality is low for an intent θ , then the person is close to a good trajectory for that intent and $\hat{\beta}_\theta$ will be high. Thus, a high $\hat{\beta}_\theta$ means that the person's input is well-explained by that intent. On the other hand, high suboptimality per time means the person is far from good trajectories, so θ 's cost model C_θ does not explain the person's trajectory and $\hat{\beta}_\theta$ will be low.

2.4 Confidence-based Arbitration

Given a confidence estimate $\hat{\beta}_\theta$ for every $\theta \in \Theta$, the robot can predict the most likely one $\theta^* = \arg \max_{\theta \in \Theta} P(\theta \mid \xi_{0 \rightarrow t}, \hat{\beta}_\theta)$ using Eq. (4). From here, one natural style of assistance is “policy blending” [8]. First the robot computes an optimal trajectory under the most likely intent, $\xi^* = \arg \min_{\xi} \sum_{x \in \xi} C_\theta^*(x)$, the first action of which is u^* . Then the robot combines u^* and $\mathcal{T}(a^t)$ using a blending parameter $\alpha \in [0, 1]$, resulting in the robot action $u^t = \alpha \mathcal{T}(a^t) + (1 - \alpha)u^*$. We also refer to α as the human's control authority.

Prior work proposes different ways to arbitrate between the robot and human actions by choosing α proportional to the robot's distance to the goal or to the probability of the most likely goal [8]. However, when using the probability $P(\theta^* \mid \xi)$, θ^* might look much better than the other intents, resulting in the robot wrongly assisting for θ^* . Distance-based arbitration ignores the full history of the user's input and can only accommodate simple intents. Instead, we propose that the robot should use its confidence in the most likely intent, $\hat{\beta}_{\theta^*}$, estimated according to Sec. 2.3, to control the strength of its arbitration:

$$u^t = \min(1, 1/\hat{\beta}_{\theta^*}) \mathcal{T}(a^t) + (1 - \min(1, 1/\hat{\beta}_{\theta^*}))u^* \quad (8)$$

When $\hat{\beta}_{\theta^*}$ is high, i.e. the robot is confident that the predicted intent θ^* can explain the person's input, α is low, giving the robot more influence through its action u^* . When $\hat{\beta}_{\theta^*}$ is low, i.e. not even the most likely intent explains the person's input, α increases, giving the person's action a^t more authority.

2.5 Using Confidence for Lifelong Learning

Estimating $\hat{\beta}_\theta$ offers the robot great flexibility in choosing how strongly to assist the user. It also gives the robot a way to detect when its set Θ is *misspecified*: if all estimated $\hat{\beta}_\theta$ for $\theta \in \Theta$ are below a threshold ϵ , the robot is missing the person's intent. Once the robot has identified that its intent set is misspecified, it should ask the person to teach it. We represent the missing intent θ_ϕ as a neural network cost parameterized by ϕ and learn it via deep maximum entropy IRL [10] (detailed in Appendix A).

Once we have a new intent θ_ϕ , the robot updates its intent set $\Theta \leftarrow \Theta \cup \theta_\phi$. The next time the person needs assistance, the robot

can perform confidence estimation, goal inference, and arbitration as before, using the new library of intents. Learned rewards fit naturally into our framework, allowing for a simple way to compare against the known intents. However, one could imagine adapting our method to the many other ways to learn an intent, from imitation learning [14, 25], to dynamic movement primitives [22]. For instance, if we parameterize intents via policies, we can derive a similar confidence metric based on probabilities of observed human actions under a stochastic policy, rather than costs.

3 EXPERT CASE STUDY

In this section, we introduce three manipulation tasks and use expert data to analyze confidence estimation and assistance. We later put CASA's assistive capacity to test with non-experts in a user study in Appendix C.

3.1 Experimental Setting

We conduct our experiments on the simulated 7-DoF JACO arm shown in Fig. 2. We use the pybullet interface [6] and teleoperate the robot via keyboard commands. We map 6 keys to bi-directional xyz movements of the robot's end-effector, and 2 keys for rotating it in both directions.

We test CASA on 3 different tasks. In the first, deemed *Known Goal*, we control for the situation where there is no misspecification: the robot must assist the user to move to the known green goal location in Fig. 2. In the other tasks, we test CASA's efficacy in the case of misspecification, where the user's desired intent is initially missing from the robot's known set Θ . In the second task, *Unknown Goal*, the person teleoperates the robot to the red goal which is unknown by the robot. Finally, in the third task, *Unknown Skill*, the person tries a more complicated intent which involves pouring the contents of the cup at a goal location unknown to the robot.

For the Unknown Goal and Unknown Skill tasks, we first run CASA before being exposed to the new intent (*CASA before learning*). Detecting low confidence, the robot asks for demonstrations and learns the missing intents via deep maximum entropy IRL as discussed in Sec. 2.5. We teleoperate with CASA *after learning*, to assess the quality of robot assistance after learning the new intent.

We compare CASA to a policy blending assistance (PBA) baseline [8] most closely related to our work that assumes $\beta = 1$ for all intents and arbitrates according to the distance d_{θ^*} to the predicted goal: $\alpha = \min(1, d_{\theta^*}/D)$, with D some threshold past which the robot does not assist. Our choice of baseline was informed by an analysis of arbitration methods presented in Appendix B.

3.2 Well-specified Tasks

Fig. 2 (top) showcases the results of our experiment for the Known Goal task. Looking at the confidence plot, we see that $\hat{\beta}_\theta$ increases with time for the correct green goal, while it remains low for the alternate known purple goal. In the arbitration plot, as $\hat{\beta}_{\theta^*}$ increases, α gradually decreases, reflecting that the robot takes more control authority only as it becomes more confident that the person's intent is indeed θ^* . Similarly, since there is no misspecification, PBA arbitration steadily decreases the human's contribution to the final control. Both methods result in smooth trajectories which go to the correct goal location.

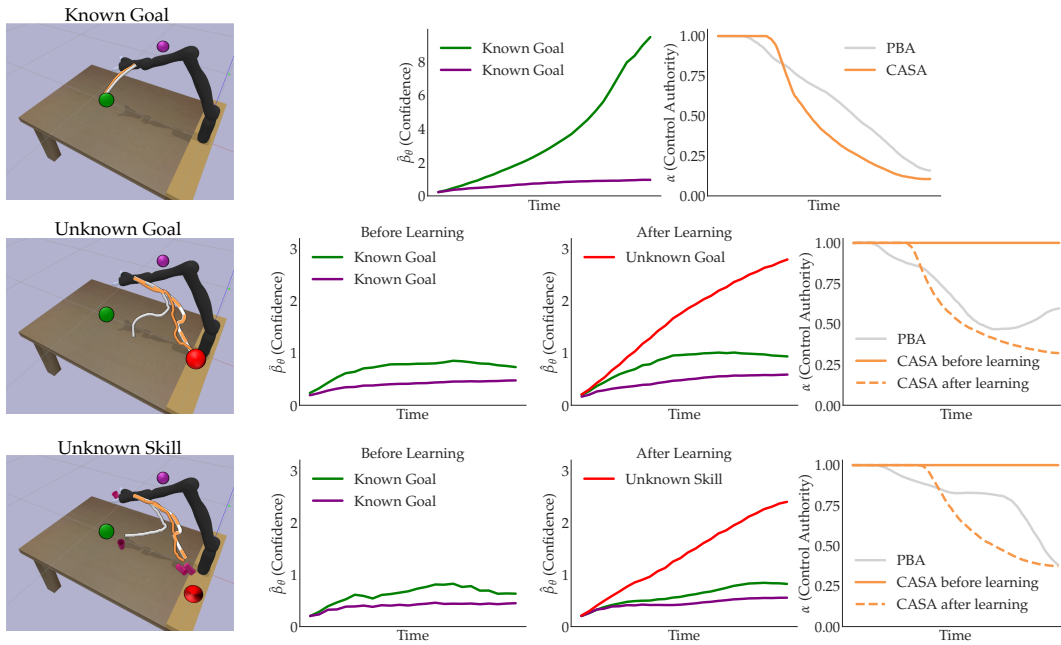


Figure 2: Expert case study results. For each of three manipulation tasks, we compute confidence estimates before learning and, for the misspecified tasks (middle, bottom), we recompute the confidence estimates after learning. We also plot the strength of assistance before and after learning and compare to a policy blending baseline [8].

3.3 Misspecified Tasks

Our approach distinguishes itself in how it handles misspecified tasks. During the Unknown Goal task, in Fig. 2 (middle), CASA before learning estimates low $\hat{\beta}_\theta$ for both goals, since neither goal explains the person’s motion moving towards the red goal. The estimated $\hat{\beta}_\theta$ is slightly higher for the green goal than for the purple one because it is closer to the user’s input; however, neither are high enough to warrant an arbitration α below 1, and thus the robot receives no control. In Fig. 2 (bottom), we observe almost identical behavior before learning for the Unknown Skill task: the known intents do not match the user’s behavior, and thus the user is given full control authority and completes the task.

This contrasts PBA, which, for both Unknown Goal and Unknown Skill, predicts the green goal as the intent. Since in both cases the user’s desired trajectory passes near the green goal, PBA erroneously takes control and moves the user towards it, requiring the human to counteract the robot’s controls to perform the task.

In the middle plots for each of the misspecified tasks, we see for CASA after learning, the newly-learned intents receive confidence estimates which increase as the robot is able to observe the user, and CASA contributes more to the control as it becomes confident.

4 USER STUDY

We tested our method’s assistive capacity with non-experts in a user study. We replicated our lab set-up in a pybullet simulator [6] in which users could teleoperate a 7 DoF JACO robotic arm using keyboard inputs (Fig. 2). We tested the same three tasks from Sec. 3 (Known Goal, Unknown Goal, and Unknown Skill) and manipulated the *assistance method* with three levels: no assistance (NA), policy blending assistance (PBA) [8], and CASA. For Unknown Goal and

Unknown Skill, we compared our method before and after learning new intents against the NA and PBA baselines.

We found that when there was no misspecification (Known Goal), CASA performed similarly to PBA, and both did better than NA. Meanwhile, when there was misspecification (Unknown Goal and Unknown Skill), assisting with CASA before learning quickly relinquished control to the human and performed similarly to NA, whereas PBA forced the users to counteract the robot’s control. CASA after learning, however, showed a significant improvement in assistance quality. Full results are described in Appendix C.

5 CONCLUSION

We formalized a confidence-aware shared autonomy process where the robot can adjust its assistance based on how confident it is in its prediction of the human intent. We introduced an approximate solution for estimating this confidence, and demonstrated its effectiveness in adjusting arbitration when the robot’s skill set is misspecified and enabling continual learning of new skills.

Due to COVID, we ran our experiments in a simulator, which does not replicate the difficulty inherent in teleoperating a real manipulator via a joystick interface. Despite this, we are encouraged to see robots have a more principled and robust way to arbitrate shared autonomy, as well as decide when they need to learn more to be better teammates. We look forward to applications of our confidence-based ideas beyond manipulation robots, to semi-autonomous vehicles or any other shared autonomy scenarios.

ACKNOWLEDGMENTS

Research supported by the Air Force Office of Scientific Research, the Office of Naval Research, and NSF grant IIS1734633.

REFERENCES

- [1] F. Abi-Farraj, C. Pacchierotti, and P. R. Giordano. 2018. User Evaluation of a Haptic-Enabled Shared-Control Approach for Robotic Telemanipulation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1–9. <https://doi.org/10.1109/IROS.2018.8594030>
- [2] Peter Aigner and Brennan McCarragher. 1997. Human integration into robot control utilising potential fields. In *Proceedings of International Conference on Robotics and Automation*, Vol. 1. IEEE, 291–296.
- [3] Chris Baker, Joshua B Tenenbaum, and Rebecca R Saxe. 2007. Goal inference as inverse planning. (01 2007).
- [4] Andreea Bobu, Andrea Bajcsy, Jaime F Fisac, Sampada Deglurkar, and Anca D Dragan. 2020. Quantifying Hypothesis Space Misspecification in Learning From Human–Robot Demonstrations and Physical Corrections. *IEEE Transactions on Robotics* 36, 3 (2020), 835–854.
- [5] Daniel S Brown, Shin-Young Jung, and Michael A Goodrich. 2014. Balancing human and inter-agent influences for shared control of bio-inspired collectives. In *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 4123–4128.
- [6] Erwin Coumans and Yunfei Bai. 2016–2019. PyBullet, a Python module for physics simulation for games, robotics and machine learning. <http://pybullet.org>.
- [7] Jacob W Crandall, Nathan Anderson, Chace Ashcraft, John Grosh, Jonah Henderson, Joshua McClellan, Aadesh Neupane, and Michael A Goodrich. 2017. Human-swarm interaction as shared control: Achieving flexible fault-tolerant systems. In *International Conference on Engineering Psychology and Cognitive Ergonomics*. Springer, 266–284.
- [8] Anca D Dragan and Siddhartha S Srinivasa. 2013. A policy-blending formalism for shared control. *The International Journal of Robotics Research* 32, 7 (2013), 790–805. <https://doi.org/10.1177/0278364913490324> arXiv:<https://doi.org/10.1177/0278364913490324>
- [9] Ahmetcan Erdogan and Brenna D Argall. 2017. The effect of robotic wheelchair control paradigm and interface on user performance, effort and preference: an experimental assessment. *Robotics and Autonomous Systems* 94 (2017), 282–297.
- [10] Chelsea Finn, Sergey Levine, and Pieter Abbeel. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*. 49–58.
- [11] Jaime F Fisac, Andrea Bajcsy, Sylvia L Herbert, David Fridovich-Keil, Steven Wang, Claire J Tomlin, and Anca D Dragan. 2018. Probabilistically Safe Robot Planning with Confidence-Based Human Predictions. *Robotics: Science and Systems (RSS)* (2018).
- [12] David Fridovich-Keil, Andrea Bajcsy, Jaime F. Fisac, Sylvia L. Herbert, Steven Wang, Anca D. Dragan, and Claire J. Tomlin. 2019. Confidence-aware motion prediction for real-time collision avoidance. *International Journal of Robotics Research* (2019).
- [13] Ray C Goertz. 1963. Manipulators used for handling radioactive materials. *Human factors in technology* (1963), 425–443.
- [14] Jonathan Ho and Stefano Ermon. 2016. Generative Adversarial Imitation Learning. In *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Eds.). Curran Associates, Inc., 4565–4573. <http://papers.nips.cc/paper/6391-generative-adversarial-imitation-learning.pdf>
- [15] Shervin Javdani, Siddhartha S Srinivasa, and J Andrew Bagnell. 2015. Shared autonomy via hindsight optimization. *Robotics science and systems: online proceedings 2015* (2015).
- [16] Emmanuel Lesaffre. 2008. Superiority, equivalence, and non-inferiority trials. *Bulletin of the NYU hospital for joint diseases* 66, 2 (2008), 150–154. <http://europepmc.org/abstract/MED/18537788>
- [17] Qinan Li, Weidong Chen, and Jingchuan Wang. 2011. Dynamic shared control for human-wheelchair cooperation. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 4278–4283.
- [18] Dylan P. Losey, Craig G. McDonald, Edoardo Battaglia, and Marcia K. O’Malley. 2018. A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human–Robot Interaction. *Applied Mechanics Reviews* 70, 1 (02 2018). <https://doi.org/10.1115/1.4039145> arXiv:https://asmedigitalcollection.asme.org/appliedmechanicsreviews/article-pdf/70/1/010804/5964415/amr_070_01_010804.pdf 010804.
- [19] D. P. Losey, K. Srinivasan, A. Mandlekar, A. Garg, and D. Sadigh. 2020. Controlling Assistive Robots with Learned Latent Actions. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 378–384. <https://doi.org/10.1109/ICRA40945.2020.9197197>
- [20] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John E Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew B Schwartz, Jennifer L Collinger, and J Andrew Bagnell. 2017. Autonomy infused teleoperation with application to brain computer interface controlled manipulation. *Autonomous Robots* 41, 6 (2017), 1401–1422.
- [21] Andrew Ng and Stuart Russell. 2000. Algorithms for inverse reinforcement learning. *International Conference on Machine Learning (ICML)* 0 (2000), 663–670. <https://doi.org/10.2460/ajvr.67.2.323>
- [22] Alexandros Paraschos, Christian Daniel, Jan R Peters, and Gerhard Neumann. 2013. Probabilistic Movement Primitives. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2616–2624. <http://papers.nips.cc/paper/5177-probabilistic-movement-primitives.pdf>
- [23] Claudia Pérez-D’Arpino and Julie A Shah. 2015. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, 6175–6182.
- [24] Siddharth Reddy, Anca D Dragan, and Sergey Levine. 2018. Shared autonomy via deep reinforcement learning. *arXiv preprint arXiv:1802.01744* (2018).
- [25] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. 2020. SQL: Imitation Learning via Reinforcement Learning with Sparse Rewards. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net. <https://openreview.net/forum?id=S1xKd24tWb>
- [26] John Schulman, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, and Pieter Abbeel. [n.d.]. Finding Locally Optimal, Collision-Free Trajectories with Sequential Convex Optimization.
- [27] John Von Neumann and Oskar Morgenstern. 1945. *Theory of games and economic behavior*. Princeton University Press Princeton, NJ.
- [28] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 3 (Chicago, Illinois) (AAAI’08)*. AAAI Press, 1433–1438. <http://dl.acm.org/citation.cfm?id=1620270.1620297>




Figure 3: Analysis of arbitration methods. After tracking an optimal trajectory for the Unknown Goal task, we show the robot’s belief and confidence estimates for each known goal (left), as well as the α values under the distance, belief, and confidence-based arbitration schemes (right).

A DEEP MAXIMUM ENTROPY IRL

We represent the missing intent θ_ϕ as a neural network cost parameterized by ϕ and learn it via deep maximum entropy IRL [10]. The gradient of the IRL objective with respect to the cost parameters ϕ can be estimated by:

$$\nabla_\phi \mathcal{L} \approx \frac{1}{|\mathcal{D}^*|} \sum_{\tau \in \mathcal{D}^*} \nabla_\phi C_\phi(\tau) - \frac{1}{|\mathcal{D}^\phi|} \sum_{\tau \in \mathcal{D}^\phi} \nabla_\phi C_\phi(\tau)$$

\mathcal{D}^* are (noisy) demonstrations of the person executing the desired missing intent via direct teleoperation, and \mathcal{D}^ϕ are trajectories sampled from the C_ϕ induced near the optimal policy, which we obtain using TrajOpt [26]. With just a few (noisy) demonstrations \mathcal{D}^* of the person executing the desired missing skill via direct teleoperation, the robot can learn an expressive cost function.

B EXPERT CASE STUDY

B.1 Arbitration Method Comparison

We compare CASA to a policy blending assistance (PBA) baseline [8] most closely related to our work that assumes $\beta = 1$ for all intents and arbitrates according to the distance d_{θ^*} to the predicted goal: $\alpha = \min(1, d_{\theta^*}/D)$, with D some threshold past which the robot does not assist. Although there exist seemingly more sophisticated arbitration schemes which are functions of $P(\theta^* | \xi)$ or the entire distribution $P(\theta | \xi)$, they are significantly less robust to task misspecification. This is because when the user teleoperates according to an unknown intent, $P(\xi | \theta)$ will be low for all known $\theta \in \Theta$, but forming $P(\theta | \xi)$ requires normalizing over all known intents, after which $P(\theta^* | \xi)$ can still be very high unless the user happened to operate in a way that appears equally unlikely under the known intents.

We analyzed this phenomenon by tracking a reference trajectory for the Unknown Goal task which moves optimally towards the unknown goal (see Fig. 2 for the task layout). We compared the performances of the distance and confidence arbitration methods, as well as a belief-based method which sets $\alpha = (P(\theta^* | \xi)|\Theta| - 1)/(|\Theta| - 1)$ (chosen so that $\alpha = 0$ when $P(\theta^* | \xi) = 1/|\Theta|$, $\alpha = 1$ when $P(\theta^* | \xi) = 1$). As shown in Fig. 3, the confidence in each goal stays low enough that the robot would have left the user in full control, while the relatively higher likelihood of one goal causes the belief $P(\theta^* | \xi)$ to quickly go to 1 and thus set the user’s control authority to 0 under the belief-based arbitration scheme.

We examined one belief-based arbitration method here, but since $P(\theta^* | \xi)$ rapidly goes to 1, any other arbitration that is a function of the belief $P(\theta | \xi)$ would similarly try to assist for the wrong goal, motivating our choice of the simpler but more robust distance-based arbitration baseline.

C USER STUDY

We now present the results of our user study, testing how well our method can assist non-expert users.

C.1 Experimental Design

Due to the COVID-19 pandemic, we were unable to perform an in-person user study with a physical robot. Instead, as described in Sec. 3, we replicated our lab set-up in a pybullet simulator [6] in which users can teleoperate a 7 DoF JACO robotic arm using keyboard inputs (Fig. 2).

We split the study into four phases: (1) familiarization, (2) no misspecification, (3) misspecification before learning, and (4) misspecification after learning. First, we introduce the user to the simulation interface by asking them to perform a familiarization task. In the next phase, we tested the Known Goal task. In the third phase, we tested the two misspecified tasks, Unknown Goal and Unknown Skill, then asked participants to provide 5 demonstrations for each intent. Finally, in the fourth phase, we retested the misspecified tasks using cost functions learned from the demonstrations.

Independent Variables: For each experiment, we manipulate the *assistance method* with three levels: no assistance (NA), policy blending assistance (PBA) [8], and Confidence-Aware Shared Autonomy (CASA). For Unknown Goal and Unknown Skill, we compared our method before and after learning new intents against the NA and PBA baselines.

Dependent Measures: Before each task, we displayed an exemplary reference trajectory to help participants understand their objective. As such, for our objective metrics, we measured *Error* as the sum of squared differences between the intended and executed trajectories, *Efficiency Cost* as the sum of squared velocities across the executed trajectory, and *Effort* as the number of keys pressed. To assess the users’ interaction experience, we administered a subjective 7-point Likert scale survey, asking the participants three questions: (1) if they felt the robot understood how they wanted the task done, (2) if the robot made the interaction more effortless, and (3) if the assistance provided was useful.

Participants: We used a within-subjects design and counterbalanced the order of the assistance methods. We recruited 11 users (10 male, aged 20-30) from the campus community, most of whom had technical background.

Hypotheses:

H1: If there is no misspecification, assisting with CASA is not inferior to assisting with PBA, and is superior to NA.

H2: If there is misspecification, assisting with CASA before learning is more accurate, efficient, and effortless than with PBA and not inferior to NA.

H3: If there is misspecification, assisting with CASA after learning is more accurate, efficient, and effortless than NA.

H4: If there is misspecification, participants will believe the robot understood what they want, feel less interaction effort, and find




Figure 4: Our user study objective metrics. For every task, we measured error with respect to an intended trajectory (left), smoothness of the executed trajectory (middle), and effort relative to direct teleoperation (right).




Figure 5: Subjective user study results. When there is no misspecification (left), our method is not inferior to PBA, whereas when there is misspecification (center, right), the participants prefer our method after learning a new intent.

the assistance more useful with CASA after learning than with any other baseline.

C.2 Analysis

Objective. Fig. 4 summarizes our main findings. For Known Goal, which is well-specified, CASA does no worse than PBA and better than NA in terms of relative effort and error. We confirmed this by running an ANOVA, finding a significant main effect for the method ($F(2, 30) = 104.93, p < .0001$ for effort; $F(2, 30) = 8.93, p = .0009$ for error). In post-hoc testing, a Tukey HSD test revealed that CASA is significantly better than NA ($p < .0001$ for effort, $p = .0013$ for error). We also performed a non-inferiority test [16], and obtained that CASA is non-inferior to PBA within a margin of 0.065 for effort, 0.025 for efficiency, and 0.26 for error. These findings are in line with H1 and were expected, since the robot should have no problem handling known intents.

For the two misspecified tasks, we first ran an ANOVA with the method (CASA before learning, NA, and PBA) as a factor, and the task as a covariate, and found a significant main effect ($F(2, 62) = 11.8255, p < .0001$ for effort; $F(2, 62) = 6.119, p = .0038$ for error). A Tukey HSD revealed that CASA is significantly better than PBA ($p = .0005$ for effort, $p = .005$ for error). We also ran a non-inferiority test, and obtained that CASA is non-inferior to NA within a margin of 0.035 for effort, 0.02 for efficiency, and 1.4 for error for Unknown Goal, and 0.03 for effort, 0.09 for efficiency, and 4.5 for error for Unknown Skill. For both unknown tasks, CASA before learning is essentially indistinguishable from NA since a low $\hat{\beta}_{\theta^*}$ would make the robot rely on direct teleoperation. Both the figure and our statistical tests confirm H2, which speaks for the consequences of confidently assisting for the wrong intent.

For efficiency cost, we did not find an effect, possibly because Fig. 4 shows that PBA is more efficient for the Unknown Skill task than other methods. Anecdotally, PBA forced users to an incorrect goal thus preventing them from pouring, which explains the lower efficiency cost. By having a high arbitration for the wrong intent, PBA can cause a smooth trajectory, since it lowers the control authority of the possibly-noisy human inputs. However, this trajectory does not accomplish the task. When running an ANOVA for each of the tasks separately, we found a significant main effect for the method for Unknown Goal ($F(2, 30) = 9.66, p = .0006$), and a post-hoc Tukey HSD revealed CASA is significantly better than PBA ($p = .0032$), further confirming H2.

Lastly, we looked at the performance with CASA after learning the new intents. For Unknown Goal, a simple task, the figure shows that CASA after learning doesn't improve efficiency and error, but it does reduce relative effort when compared to NA. For Unknown Skill, a more complex task, CASA after learning outperforms NA. This is confirmed by an ANOVA with the method (NA, CASA after learning) as the factor, where we found a significant main effect ($F(1, 41) = 53.60, p < .0001$ for effort; $F(1, 641) = 8.6184, p = .0054$ for efficiency cost), supporting H3.

Subjective. We show the average Likert survey scores for each task in Fig. 5. In line with H1, for the Known Goal task, users thought the robot under both PBA and CASA had a good understanding of how they wanted the task to be done, made the interaction more effortless, and provided useful assistance. The results are in stark contrast to NA, which scores low on all those metrics. For Unknown Goal and Unknown Skill, all methods fare poorly on all questions except for CASA after learning, supporting our H4.